

Parallel I/O on the IBM Blue Gene/L System

Abstract

Blue Gene/L is currently the world's highest performing and most scalable parallel system. With a raw compute power varying from 5.7 (single-rack configuration) to almost 370 TF (64-rack configuration), it can enable scientific and engineering computations that were just not possible before. A system with this much compute power can be very demanding on its I/O subsystem, however. Many important applications require the processing of vast volumes of data. With up to 1 TB of main memory per rack, checkpoint/restart can also require the saving and restoring of a lot of data. This article focuses on the issue of supporting high-performance file I/O on the Blue Gene/L system.

1 Introduction

Blue Gene/L is an exciting new machine. With its unprecedented amounts of computing power and scalability, it enables scientific and engineering computations that have never been possible before. The Blue Gene/L system is different from most current parallel machines in that the compute nodes run a very lightweight kernel, called the compute node kernel (CNK). One feature of the CNK is that it forwards system calls off the compute node for service. A separate process, running on the I/O node, services system calls on behalf of the CNK.

In this article we discuss how parallel I/O is implemented on the Blue Gene/L machine and the available options for parallel file systems.

2 Background

A high-level block diagram of a Blue Gene/L system is shown in Figure 1. The organization of a Blue Gene/L system is based on the principles of physical partitioning and specialization. The compute nodes are specialized to run user application processes. The I/O nodes are specialized to perform system functions related to job control and I/O. The data servers are specialized to store data files used by the user applications. BG/L machines may be configured with as few as one I/O node per 64 compute nodes or as many as one I/O node per 8 compute nodes.

The architecture of the data server cluster varies widely between Blue Gene system deployments, as we will see in this article. The data servers are typically implemented with conventional off-the-shelf servers. In some cases, the data servers have the disks internal to them. In other cases, the data servers are connected to disks through a storage-area network. Depending on the file system operating in the Blue Gene/L I/O nodes, the data servers can be file servers (e.g., NFS servers for an NFS file system) or just storage servers (e.g., NSD servers for a GPFS file system).

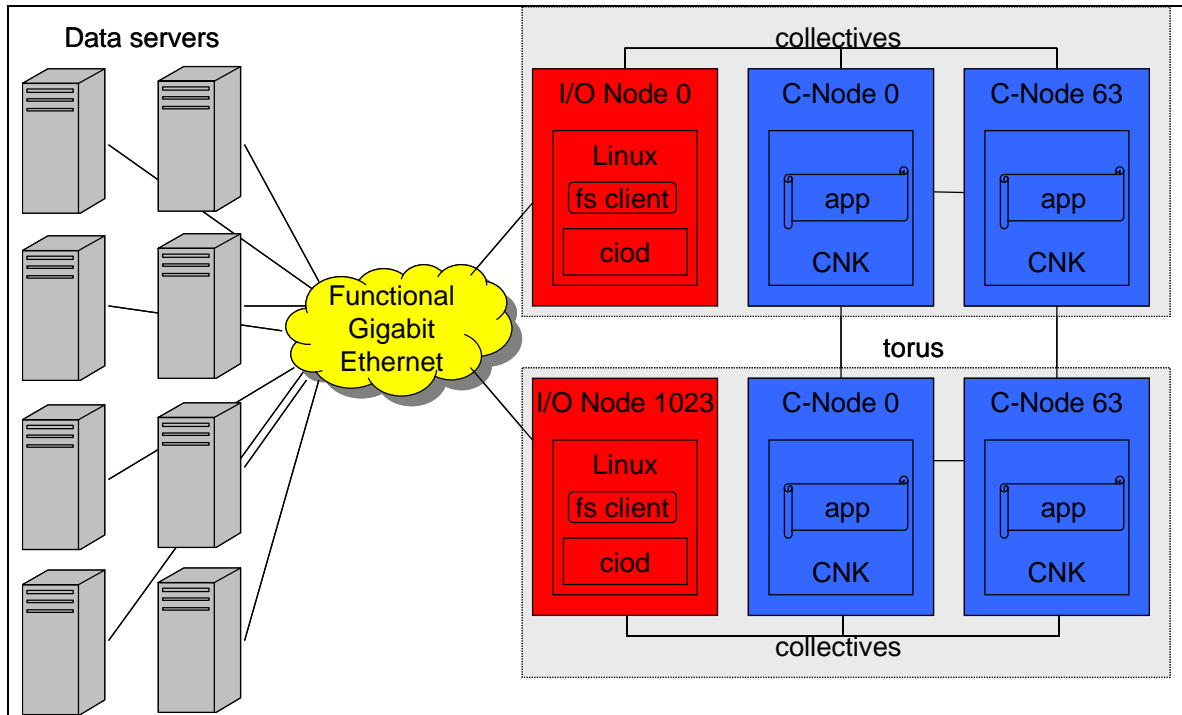


Figure 1: High-level block diagram of a Blue Gene/L system. The figure shows the compute and I/O nodes of the Blue Gene/L core and also the data servers that are connected to it.

User application processes run only on the compute nodes. Furthermore, the model is to have a single thread of execution per processor of the compute node. This can be accomplished either with a single process (in coprocessor mode) or with two independent processes (in virtual node mode) per compute node. (Note that in coprocessor mode, the second processor is used to perform message-passing functions.)

The compute node kernel was designed to implement only the strictly necessary functionality for this model of execution, thus minimizing any perturbation of the running application processes. In particular, all I/O system calls initiated by user application processes running on a compute node (e.g., open, close, read, write) are trapped by the CNK, and a request to perform the indicated function is shipped to the corresponding I/O node.

The I/O request shipped by CNK is received in the I/O node by a process called the Control and I/O Daemon (CIOD). This is a user process that runs with the same user id as the owner of the job running on the compute nodes. The CIOD simply reissues the request as a system call to the I/O node operating system.

In addition to input/output, the I/O nodes are expected to perform other system functions. We have chosen to run Linux as the operating system; hence, the I/O nodes can support any file system for which there is an appropriate port to Linux PowerPC 32-bit.

Compute and I/O nodes are interconnected through the *collectives network*, also known as the tree network. This is the same network that supports reduction and broadcast operations among the compute nodes. In the case of function shipping, this network is used as a point-to-point network between a compute node and its corresponding I/O node.

I/O operations are always initiated by an application process. The CNK will compose a message describing the request and send it to the I/O node for execution by CIOD. Messages in the collectives network must be decomposed into individual packets of 256 bytes and then reconstructed at the destination (CIOD for requests, CNK for replies). Most typical I/O operations can be implemented as single messages and many of them as single packets.

Very large file reads and writes, however, are decomposed into multiple messages, and each message is executed in order before proceeding to the next message. For example, a 5 MB read operation from the application process can be decomposed into 10 smaller read operations by the CNK. Each operation is then performed through a message that goes to CIOD for execution and then comes back with the resulting data. The CNK will complete all individual messages before finishing the read request and returning control to the application process.

The maximum size of messages between compute and I/O nodes can be controlled by the administrator for a particular installation through system configuration files. The CIOD environment variable `CIOD_RDWR_BUFFER_SIZE` controls the value of the buffer size for `read` and `write` system calls. The default value is 87600 bytes. “Popular” values that have shown good performance in installed systems are 256 KB and 512 KB, although some groups use values as large as 2–4 MB. The CIOD environment variables are summarized in Table 1.

CIOD supports a service connection that provides a trace facility and status reporting. The intent is to let a system programmer or administrator check on a running CIOD without disturbing it. CIOD starts a secondary thread that monitors the service connection so there is minimal impact to the work being done by the main thread. CIOD uses TCP port 7201 for the service connection. The trace facility supports turning tracing on and off for categories of events. By default, the trace output is sent back across the service connection but can also be directed to a file. The trace output allows one to see the flow of control through a running CIOD. The status reporting allows one to see the current state of various objects maintained by CIOD. One can check on individual compute nodes and get overview performance information for certain system calls.

Table 1: CIOD environment variables.

Variable name	Description
CIOD_RDWR_BUFFER_SIZE	The value is the buffer size for read and write system calls. The value can be increased or decreased to find the best match for the network and file servers used by the Blue Gene system. CIOD allocates one buffer of the specified size for each compute node that it is managing. A larger value causes CIOD to allocate more memory. The default value is 87600.
CIOD_TREE_RECV_RETRY	The value is the number of times CIOD tries to receive a packet from the tree device before deciding a packet is not available and yielding the processor. A value of 16 causes CIOD to never yield and return immediately. A larger value causes CIOD to pay more attention to the tree. The default value is 16.
CIOD_TREE_SEND_RETRY	The value is the number of times CIOD tries to send a packet to the tree device before yielding the processor. A value of 0 causes CIOD to never yield and loop until the packet is sent. A larger value causes CIOD to pay more attention to the tree. The default value is 0.
CIOD_TREE_MULTIPLIER	The value is the default multiplier controlling how often CIOD checks for messages from the compute nodes before checking the control connections and waiting I/O. The tree multiplier is dynamically adjusted by CIOD when at least one compute node is using a debugger or there is waiting I/O. A larger value causes CIOD to pay more attention to messages from compute nodes. The default value is 8192.
DEBUG_SOCKET_STARTUP	The value is a string that controls what blocks have the CIOD service connection enabled. The value ALL enables the service connection on every block. Otherwise, the value is the name of one block where the service connection is enabled. When the variable is not set, the service connection is not enabled.

As previously noted, the I/O nodes can use any file system that has been ported to their version of Linux for PowerPC 32-bit. As part of the boot process, every I/O node mounts from a primary NFS server that contains the code for the rest of the boot. Once that primary file system is mounted, the I/O nodes can mount other file systems. The rest of this article discusses the options for parallel file systems that are available for Blue Gene/L.

3 File Systems

Currently four file system options are being actively pursued for Blue Gene/L systems: NFS, GPFS, PVFS2, and Lustre.

3.1 NFS

The first option is to use NFS to mount a file system. This works with almost any file system, parallel or not. If a parallel file system is available, it may be NFS-exported from a variety of data storage machines, and each I/O node can NFS-mount from a different machine. This is a cheap and easy way to deploy a parallel file system for Blue Gene/L. Scalability limitations are associated with this configuration, however, since NFS is not really aware of the parallel nature of the underlying file system. In experiments with NFS-exported GPFS, the performance is quite good when each application process accesses a different file. On the other hand, the performance is very bad when multiple application processes share the same file.

3.2 GPFS

The second option is to use GPFS, the parallel file system developed and supported by IBM. The San Diego Supercomputer Center (SDSC) has been running a beta version of GPFS since April 2005. In December 2005, IBM released the first official version of GPFS for Blue Gene.

The most common configuration for a cluster running GPFS has the disks attached via so-called Network Shared Disk (NSD) servers. This configuration allows operation over an IP networking infrastructure and has been shown to scale well. In a conventional cluster the compute nodes function as GPFS clients, whereas for Blue Gene the I/O nodes constitute the cluster and function as the clients. The compute nodes, which run only a lightweight kernel, use the tree network to move file I/O to and from the associated I/O nodes. SDSC's single-rack Blue Gene has 128 I/O nodes, the maximum allowable. This I/O-rich configuration enables high I/O performance.

Two specific parallel file systems are provided: bggpfs and gpfs-wan. Both are accessible via GPFS and use IA-64 nodes as the NSD servers. The bggpfs system is dedicated to Blue Gene. It has 12 NSD data servers, 2 NSD metadata servers, and 20 TB of disk. By contrast, gpfs-wan is shared with other supercomputers at SDSC as well as across the TeraGrid. The gpfs-wan system is much larger than bggpfs, with 58 data servers, 6 metadata servers, and 220 TB of disk. In each case the servers are attached via Gigabit Ethernet to the I/O nodes and via Fibre Channel to the disk. Between the I/O nodes and the NSD servers is a Gigabit Ethernet switch with a single link to each server. These links limit the peak performance possible with each file system.

Performance

I/O performance of both file systems has been measured with the IOR benchmark code from LLNL. Figure 2 shows the results of scaling scans obtained running in coprocessor mode with the default node mapping. Also shown are the peak rates possible with each file system. The peak rate curves are the same for both file systems up to 32 processors. For 64 processors and higher, the peak rate using bggpfs is at a plateau of 1.5 GB/s set by the links to the 12 data servers. The peak rate curve using gpfs-wan increases beyond 32 processors until it reaches a plateau of 7.25 GB/s set by the links to the 58 data servers.

The steps in the peak rate curves arise when the ratio of compute nodes to I/O nodes increases for the default mapping. That ratio increases from 1 to 2 going from 1 processor to 2 processors, from 2 to 4 going from 8 processors to 16 processors, and from 4 to 8 going from 64 processors to 128 processors. The write rate curves for both file systems are between 50% and 65% of the corresponding peak rate curves over the entire processor range except for 1024 processors with gpfs-wan. The behavior of the read rate curves is more complicated.

For gpfs-wan the write rate reaches a maximum of just over 4.0 GB/s on 512 processors and then drops somewhat to 3.2 GB/s on 1024 processors. By contrast, the read rate reaches a maximum of 2.2 GB/s on 1024 processors after a modest drop from 256 processors to 512 processors. For bggpfs the write and read rates are the same and constant at about 800 MB/s for 64 processors and above. For gpfs-wan at higher processor counts the write and read rates vary considerably from one IOR run. The reason is probably that the file system is shared with other systems. The rates shown are the best from several runs. By comparison, the write and read rates for bggpfs vary relatively little from run to run.

The results discussed so far were obtained in coprocessor mode because that simplifies making scaling scans. Most users, however, run in virtual node mode, so additional results were obtained in that mode at the higher processor counts. For gpfs-wan the write and read rates on 512 processors and 1024 processors in virtual node mode are lower than those in coprocessor mode and reach final values of 3.1 GB/s for writes and 1.6 GB/s for reads on 2048 processors. For bggpfs the write and read rates on 512 processors and 1024 processors in virtual node mode are essentially the same as those in coprocessor mode and are sustained on 2048 processors.

In summary, rates of 4.0 GB/s for writes and 2.2 GB/s for reads have been achieved during shared use of gpfs-wan in coprocessor mode, while rates of about 75% of those have been achieved in virtual node mode.

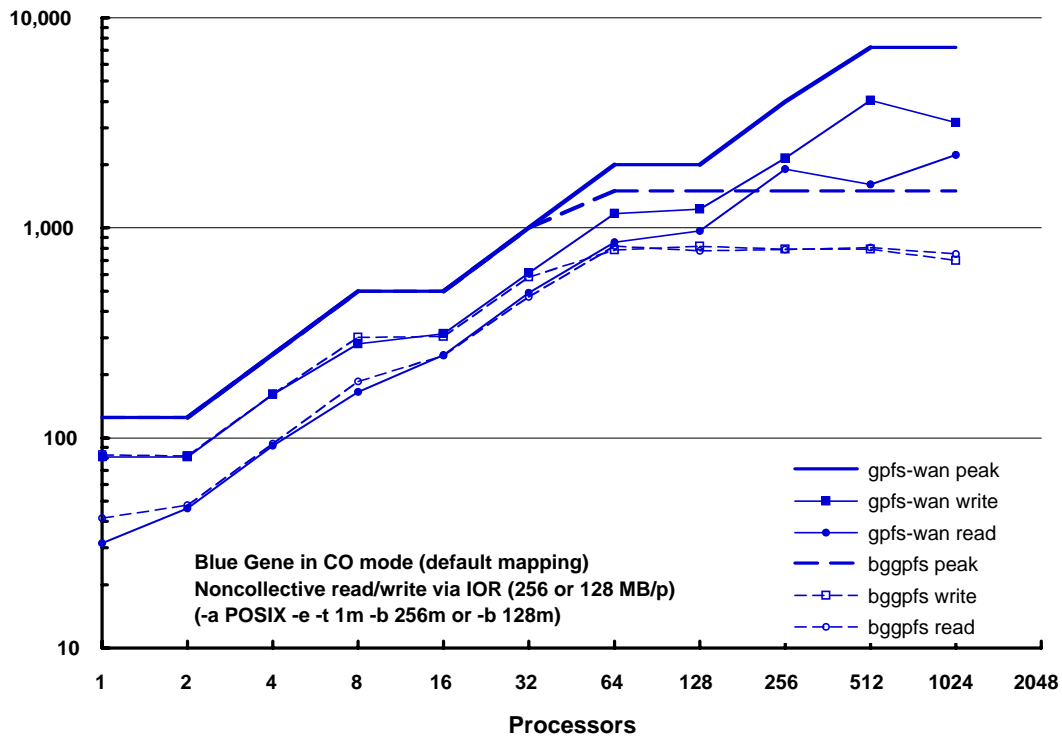


Figure 2: Benchmark results from BG/L system at SDSC running GPFS.

3.3 PVFS2

The third option is PVFS2. PVFS2 is an open source parallel file system developed by researchers at Argonne, Clemson University, and Ohio State Supercomputer Center, with help from a broad community of developers and users. While not officially supported by IBM, replacement I/O node images are available to run PVFS2 on BG/L as part of the ZeptoOS project. PVFS2 is currently installed at Argonne and at MIT, and it has been tested on the BGW system at IBM Watson.

The PVFS2 system uses a set of PVFS2 servers to allow PVFS2 clients to access the shared file system space. On BG/L, I/O nodes are PVFS2 clients. The ZeptoOS I/O node image provides the kernel module and user-space daemon that PVFS2 clients use to contact PVFS2 servers. BG/L compute nodes use the ciod system to forward I/O operations to I/O nodes, just as with the other file system options. BG/L login nodes also mount the file system, so that users can access data stored on the PVFS2 file system between their application runs and move data on and off the file system conveniently.

The I/O nodes use a Gigabit Ethernet network to communicate with PVFS2 servers. At Argonne, 14 PVFS2 servers combine to provide a modest 6.7 TB file system. Servers are xSeries 346 nodes, each with a 3.4 GHz Xeon processor, 4 GB of RAM, and 800 GB of 10K SCSI HDDs on a ServeRAID 6i+ controller in RAID5 mode. The built-in GigE

connections were replaced with Intel Pro 1000/MT GigE cards in order to allow for the use of jumbo frames. This configuration peaks at 1.2 GB/s read and 620 MB/s writes, and a single process can create approximately 50 files/second.

As part of the IBM BGW Consortium Day, the PVFS2 team had an opportunity to test PVFS2 on the BGW system. This system has a larger number of compute nodes and was configured with more PVFS2 servers, 33, during testing. The disk resources on these servers were very modest, significantly limiting overall write performance. Results from this testing are shown in Figure 3.

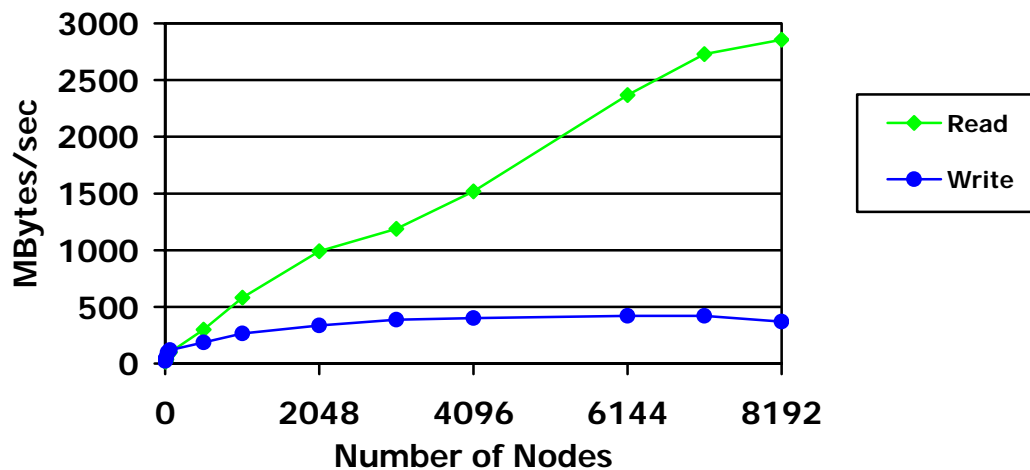


Figure 3: Benchmark results from BG/L system at IBM Watson running PVFS2.

PVFS2 is the first parallel file system to be made freely available for the BG/L system. Installation is not trivial, but an experienced system administrator should have no problem. For large I/O operations, performance is quite good, and the PVFS2 team is actively working to improve the performance of the system overall, including both small I/O operations and metadata operations, such as listing directories and creating and removing many small files.

3.4 Lustre

The fourth option is to use the Lustre file system. Lustre is a scalable cluster file system designed, developed, and maintained by [Cluster File Systems, Inc.](http://www.lustre.org/index.html) For more about CFS and Lustre in general, please visit <http://www.lustre.org/index.html>. Lawrence Livermore National Laboratory's 64K Blue Gene/L system uses Lustre as its parallel file system.

LLNL Hardware Configuration

Livermore's 64-rack BG/L system uses an I/O configuration of 64 compute nodes for every I/O node (see Figure 4). BG/L has 1024 I/O nodes. Each I/O node is a Lustre client and is attached via Gigabit Ethernet to a system of 15 switches (the federated switch and BG/L and Lustre edge switches, to be described later) that connects to the 896 TB Blue Gene Lustre Cluster (BLC) file system via 448 Gigabit Ethernet links. BLC is a cluster of 224 Intel dual Nocona nodes with 2 GB memory and 2.8 GHz processors. Each of these Nocona nodes is a Lustre "Object Storage Server" (OSS). Each OSS has 2 Gigabit Ethernet links coming into it and connects to a Data Direct Networks 8500 Controller via a single 2 Gb/s host interface. The disk storage is provided by 28 racks of Data Direct Networks (DDN) RAID storage. Each rack of DDN storage contains 32 TB of usable SATA storage accessible via eight 2 Gb/s fibre channel interfaces.

The link between BG/L and storage is provided by a network of Cisco 6509 switches. Currently the 1024 1 Gb/s ION interfaces associated with BG/L and the storage purchased with it is attached to "edge switches." The edge switches are connected to a cluster of six "core" switches via 10 Gb/s connections. Two other systems currently "share" the BLC file system by way of the federated network. The first of these is the Storage Lustre Interface Cluster (SLIC), a cluster of 10 Itanium Tiger 4 nodes that is used to transfer files out of the Lustre file system into the High Performance Storage System (HPSS) data archive for permanent storage. Users log into SLIC and use either parallel ftp or htar to move their files to HPSS. The other cluster with access to BLC is Gauss. Gauss is a high-performance visualization cluster of 256 dual opteron (6 GB memory per CPU) Infiniband interconnected nodes, using NVIDIA 4500 graphics and each with a GE link into the federated switch that connects it to the Blue Gene Lustre cluster file system. Gauss will be used exclusively by users of BG/L for data visualization.

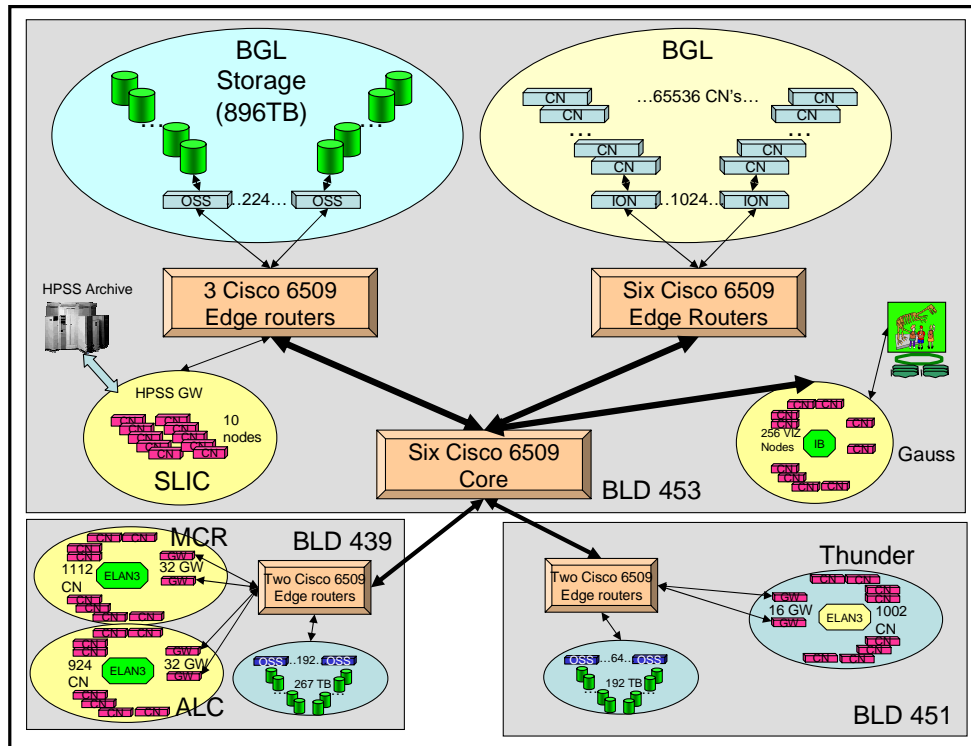


Figure 4: Lustre storage network architecture at LLNL.

Projected Performance

With its 1024 1 Gb/s interfaces the BG/L system at LLNL has a potential I/O bandwidth of 128 GB/s. The original storage design specified an I/O bandwidth of 40 GB/s—far below LLNL’s normal I/O bandwidth to CPU performance ratio. Cost necessitated scaling back this original design goal. The adjustments began with the disk subsystems. At the time the disks were purchased, there was a choice between 146 GB fibre drives for over \$1000 each or enterprise-class 250 GB SATA drives for less than \$500 each. Twice the storage for half the cost is a compelling argument. LLNL purchased 28 racks of SATA storage. Each rack has a theoretical bandwidth of 1.2 GB/s for a total of 33.6 GB/s. The trade-off made was in terms of performance. SATA disks provide very good streaming performance but poor random performance. With a ratio of 65,536 (or 131,072 in virtual node mode) clients to 448 RAID devices, the LLNL workload is random. Based on the disk specifications a request was expected to take 18 ms (\approx 55 requests per second) or, given a 1 MB block size, about 55 MB/s per RAID device. The 448 RAID devices were therefore expected to achieve around 24.8 GB/s.

The LLNL Lustre network was designed to provide 45 GB/s theoretical bandwidth between BG/L and its associated storage cluster. The current configuration uses six Cisco 6509’s as BG/L edge switches (connecting the BGL core to the federated switch infrastructure) and three 6509’s as disk edge switches (connecting the 224 file system OSS nodes to the federated switch). The whole configuration is tied together with another six 6509’s acting as core switches. The Lustre core architecture was also designed to

provide a connection point for the other parts of the LLNL Lustre-based “site wide global file system” (SWGFS). The network architecture is based on layer 3 switching and uses “equal cost multipathing” (ECMP) and link aggregation where appropriate. The design goal was to have the network deliver 80% of theoretical bandwidth, or 36 GB/s.

Actual Performance

LLNL staff first analyzed the capabilities of each I/O-related element in the path between the user application code and the physical disk devices independently. Then the staff tested groups of elements leading up to a user-level I/O benchmark (IOR) designed to simulate actual user behavior. The lessons learned and problems found and fixed are far too extensive to discuss here, but the bottom line is that, as expected, the disks have proven to be the limiting factor related to I/O performance. With exceptional cache management and carefully selected block sizes, write performance can exceed raw disk random I/O rates. The best full-system write performance achieved is around 26 GB/s. Improving read performance has proven to be a harder task, given the number of concurrent data streams per physical device, the amount of memory available in the I/O nodes for caching, and the synchronous nature of CIOD. The best read performance achieved is around 22 GB/s (not achieved in the run below). The results in Figure 5 were generated using the IOR benchmark on November 25, 2005.

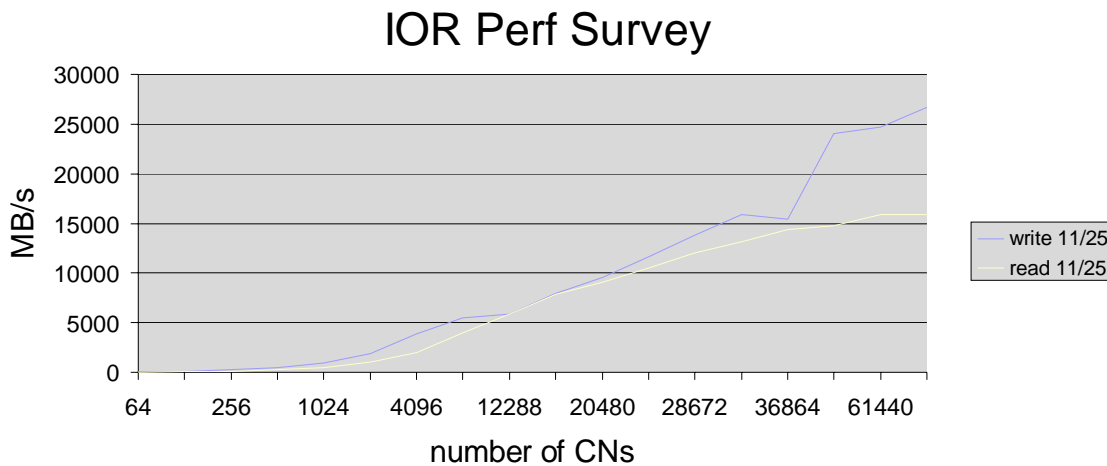


Figure 5: Benchmark results from BG/L system at LLNL running Lustre.

Summary

An enormous amount of work has been done during the past year to make Lustre perform efficiently and reliably on Livermore’s BG/L system. The approach was to isolate and tune each component and then combine and tune again. The conglomeration of components involved in I/O on BG/L, from the Lustre client and server software itself, to the tunables of the DDN controllers and the network edge and core switches, as well as the buffering in the CIOD on the ION, have all been tweaked many times in the quest to improve performance. And there is still work to do. All performance numbers listed in this study were run in coprocessor mode. Performance in virtual node mode (the second processor is used for calculation tasks as well as communication) is substantially below

coprocessor mode performance, principally because almost no time has yet been spent studying or tuning virtual node mode performance. Metadata performance is also an area that requires further study and improvement, as is simultaneous access to the Lustre file system from all systems. During 2006, LLNL will continue to study and improve remaining performance and reliability issues.

BGL Lustre Availability

Lustre is available for BGL style machines however it is not explicitly labeled as such on the Lustre website. From Lustre's point of view the BGL nodes are simply low memory ppc nodes so there is no special `_BGLonly_` version of the code. LLNL is currently running v1.4.5.8 of the Lustre code plus perhaps a dozen locally maintained BGL specific patches for mostly scalability and performance issues. Many of these patches have already been committed upstream to newer versions of Lustre. However other new features have also been added upstream that LLNL has not yet fully tested, so we're sticking with a v1.4.5.8 based version.

We suggest users wishing to run Lustre on BGL client IONs download version 1.4.5.8 from the Lustre website and compile it as follows for ppc (on a ppc node). They will of course also have to patch their kernel with the proper Lustre patches.

```
> sh autogen.sh
> ./configure --with-linux=<your-linux-source> --disable-liblustre \
  --disable-docs --disable-server --disable-readline --disable-tests
> make rpms
```

This will work well for small BGL style systems, of 1-2K nodes, beyond that scale there may be some issues encountered that LLNL is maintaining patches for. If more information or assistance with Lustre is needed please feel free to contact Kim Cupps at LLNL, cupps2@llnl.gov

4 Conclusions

Multiple options are available for attaching storage to BG/L systems. For systems with very light I/O loads, an NFS solution is probably adequate. For heavier I/O loads, a parallel file system is preferable. If commercial support is a concern, IBM's GPFS is currently the best choice. For a free, community-supported solution, PVFS2 and Lustre are excellent options and can be matched with a wide variety of storage hardware.

References

ANL BG system	http://bgservice.bgl.mcs.anl.gov/bgl/
ANL MCS	http://www-fp.mcs.anl.gov/division/welcome/default.asp
BG Consortium	http://www-fp.mcs.anl.gov/bgconsortium/
GPFS	http://www-03.ibm.com/servers/eserver/clusters/software/gpfs.html
IBM Research	http://www.research.ibm.com/bluegene/
LLNL	http://www.llnl.gov/asci/platforms/bluegene1/
Lustre	http://www.lustre.org/index.html
PVFS2	http://www.pvfs.org/pvfs2/
SDSC	http://www.sdsc.edu/
ZeptoOS	http://www-unix.mcs.anl.gov/zeptoos/